



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/763,099	01/21/2004	Charles Frank	119127-160997	7524
25943 7590 08/12/2010 Schwabe Williamson & Wyatt PACWEST CENTER, SUITE 1900 1211 SW FIFTH AVENUE PORTLAND, OR 97204				
EXAMINER DILLON, SAMUEL A				
ART UNIT 2185		PAPER NUMBER		
MAIL DATE 08/12/2010		DELIVERY MODE PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/763,099

Applicant(s)

FRANK ET AL.

Examiner

SAMUEL DILLON

Art Unit

2185

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 20 May 2010.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 17-28, 30-41 and 43-46 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 17-28, 30-41 and 43-46 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO/SB-06)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

1. Applicant's submission filed on May 20, 2010 has been entered. Per the amendment, Claims 17-23, 25-28, 30-33, 36-41 and 43 have been amended and Claims 45 and 46 have been added.

I. RESPONSE TO AMENDMENT(S) / ARGUMENT(S)

2. Applicant's arguments with respect to the 35 U.S.C. 103(a) rejections of Claims 17-28, 30-41, 43 and 44 have been fully considered and are persuasive, but are moot in view of the new ground(s) of rejection, as described below.

3. Regarding all other Claims not specifically traversed above and whose rejections were upheld, the Applicant contends that the listed claims are allowable by virtue of their dependence on other allowable claims. As this dependence is the sole rationale put forth for the allowability of said dependent claims, the Applicant is directed to the Examiner's remarks above. Additionally, any other arguments the Applicant made that were not specifically addressed in this Office Action appeared to directly rely on an argument presented elsewhere in the Applicant's response that was traversed, rendered moot or found persuasive above.

II. REJECTIONS BASED ON PRIOR ART

Claim Rejections - 35 USC ' 103 – Wang and Lee

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. **Claims 17-21, 23-28, 30, 31, 33-41, 43 and 44-46** are rejected under 35 U.S.C. 103(a) as being unpatentable over Wang et al. (*US Patent 6,834,326*) in view of Lee et al. ("*Petal: Distributed Virtual Disks*").

6. As per **Claim 28**, Wang disclose(s) a method comprising:

establishing a plurality of characteristics (*configuration information, including LUN and ID information, RAID type, what multicast group they are in, what extent configurations are on the disk, stripe size, etc, col 11 lns 51-61*) associated with an array group (*array of disks in a logical group, col 8 lns 31-41*) that includes a plurality of array partitions (*contents of disks in the logical group, col 8 lns 31-41*), the plurality of characteristics including a type of the array group, which indicates how data is distributed across the plurality of array partitions (*RAID type, col 11 lns 51-61*), and

establishing, on a storage medium, an array partition of the plurality of array partitions (*the NetSCSI autoconfiguration protocol is used to assign LUN and ID information to each disk, col 11 lns 51-61*).

Wang discloses that it is possible to multicast the data to each disk and have the disks select which data they need (*col 12 lns 3-9*). More specifically, Wang discloses

receiving, via a network interface (*connection between the RAID controller and each disk, fig 5*), a data access command multicast to the plurality of array partitions (*multicast the data to the disks, col 12 lns 3-9*), and determining that the data access command pertains to the array partition based at least in part on the plurality of characteristics (*the disks know what LUNs they contain, and therefore what data would pertain to them, col 11 lns 51-61, col 12 lns 3-9*).

At the time of the invention, it would have been obvious to a person having ordinary skill in the art to modify Wang to multicast the data to the disks and have the disks select which data they need. Wang discloses that there are pros (*col 12 lns 4-5*) and cons (*col 12 lns 5-9*) for this approach, but the Examiner asserts that the motivation would have been that it would reduce the burden on the RAID controller (*col 12 lns 4-5*). Therefore, it would have been obvious to modify Wang to multicast data to each disk regardless of relevance for the benefit of reducing the burden on the RAID controller.

Wang discloses configuring the disks so that they "know" some of the relevant configuration information (*col 11 lns 51-61*). For the purposes of this rejection, Wang is not relied upon to disclose establishing a root partition containing the plurality of characteristics on the storage medium containing the array partition.

Lee discloses a distributed storage system where information that describes the members of the system is replicated across all peer servers in the storage system (*section 2 paragraph 3*). More specifically, Lee discloses establishing a root partition (*location of the information on the Petal server, section 2 par 3*) containing a plurality of characteristics (*information describing the current members of the storage system and the currently supported virtual disks, section 2 par 3*) on the storage medium containing

Art Unit: 2185

an array partition (*disk storage, figure 2*) of a plurality of array partitions (*there are a plurality of disk storages, figure 2*).

Wang and Lee are analogous art in that they deal with networked storage devices. At the time of the invention, it would have been obvious to a person having ordinary skill in the art to modify Wang to store the complete configuration information in each of the networked storage devices in place of using a RAID controller, using the teachings of Lee. The motivation for doing so would have been that it is better able to tolerate and recover from any single component failure (*Lee, section 1, list after paragraph 2*). Therefore, it would have been obvious to modify Wang per the teachings of Lee for the benefit of fault tolerance, to obtain the invention of Claim 28.

7. As per Claim 17, Wang and Lee disclose(s) a storage appliance comprising:

a network interface (*Lee, connection between storage server and scalable network, figure 2*);

a storage medium (*Wang, disk in a logical group, col 8 lns 31-41; as modified by Lee to operate as a Petal server, section 2 and figure 2*); and

a controller (*Lee, functionality to operate as a Petal server, section 2 and figure 2*) coupled to the network interface and the storage medium and configured:

to establish a root partition (*Lee, location of the information on the Petal server, section 2 par 3*) on the storage medium, the root partition defining a plurality of characteristics (*Wang, configuration information, including LUN and ID information, RAID type, what multicast group they are in, what extent configurations are on the disk, stripe size, etc, col 11 lns 51-61; Lee, information describing the current members of the storage system and the currently supported virtual disks, section 2 par 3*) of an array group (*Wang, array of disks*

in a logical group, col 8 lns 31-41) that includes a plurality of array partitions (*Wang, contents of disks in the logical group, col 8 lns 31-41*), the plurality of characteristics to include a type of the array group, which indicates how data is distributed across the plurality of array partitions (*Wang, RAID type, col 11 lns 51-61*), or a parity rule of the array group,

to establish an array partition on the storage medium, the array partition being one of the plurality of array partitions (*Wang, the NetSCSI autoconfiguration protocol is used to assign LUN and ID information to each disk, col 11 lns 51-61*),

to receive, via the network interface, a data access command multicast to the plurality of array partitions (*Wang, multicast the data to the disks, col 12 lns 3-9*), and

to determine that the data access command pertains to the array partition based at least in part on the plurality of characteristics (*Wang, the disks know what LUNs they contain, and therefore what data would pertain to them, col 11 lns 51-61, col 12 lns 3-9*).

8. As per Claims 18 and 37, but more specifically to Claim 18, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the controller is further configured to receive, via the network interface, a plurality of partition commands from a host (*Lee, Petal RPC interface, section 3 paragraph 3*); and to establish the root partition and the array partition based at least in part on the plurality of partition commands (*Lee, section 3 paragraph 3*).

9. As per Claims 19, 30 and 38, but more specifically to Claim 19, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the plurality of characteristics

includes a multicast set associated with the array group (*Wang, what multicast group they are in, col 11 lns 51-61*).

10. As per Claim 20, Wang and Lee disclose(s) the storage appliance of claim 19, wherein the controller is configured to receive a set command from a host via the network interface, and to establish the multicast set associated with the array group based at least in part on the set command (*Wang, col 11 lns 51-61*). However, for the purposes of this rejection, Wang and Lee do not disclose the set command being multicast. Wang discloses using multicast packets when beneficial (*col 11 lns 63-65*). It would have been obvious to modify Wang and Lee to multicast the configuration commands, per the teachings of Wang. The motivation would have been that multicasting information can improve performance (*Wang, col 3 lns 24-33*). Therefore, it would have been obvious to modify Wang and Lee per the further teachings of Wang for the benefit of performance, to obtain the invention of Claim 20.

11. As per Claims 21, 31, 39 and 43, but more specifically to Claim 21, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the data access command is multicast to the plurality of array partitions using an Internet Protocol address (*Wang, column 8 lines 31-41*).

12. As per Claim 23, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the plurality of characteristics includes the type of the array group and a description of the plurality of array partitions (*Wang, col 11 lns 51-61*) and the controller is further configured to determine that the data access command pertains to the array partition based at least in part on the type of the array group and the description of the plurality of array partitions (*as combined above, Wang, col 12 lns 3-9*).

13. As per Claims 24 and 34, but more specifically to Claim 24, Wang and Lee disclose(s) the storage appliance of claim 23, wherein the type is a stripe (*Wang, RAID*

Art Unit: 2185

0, col 4 *Ins* 18-29) and the plurality of characteristics further includes a length of the stripe (*Wang*, col 11 *Ins* 51-61).

14. As per Claims 25 and 35, but more specifically to Claim 25, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the plurality of characteristics includes the parity rule of the array group (*Wang*, *RAID type includes what parity rule is in effect*, col 4 *Ins* 18-29; *RAID type is included in the configuration information*, col 11 *Ins* 51-61).

15. As per Claim 26, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the plurality of array partitions are associated with a plurality of logical block addresses (LBAs) (*Wang*, col 20 *Ins* 31-42) and the controller is further configured to calculate, based at least in part on the plurality of characteristics of the array group defined in the root partition, which LBAs of the plurality of LBAs are associated with the array partition (*Wang*, "what extent configurations are on the disk" is included in the configuration of the disk, col 11 *Ins* 51-61; as combined above, *Wang* discloses that disks know what they contain based on the configuration, col 12 *Ins* 3-9).

16. As per Claim 27, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the controller is configured to receive the data access command from a host and the controller is further configured to transmit, via the network interface, data directly to another array partition of the plurality of array partitions based at least in part on the data access command (*Lee*, when a new server is added and the data is still being moved to balance the servers, commands can be forwarded directly between servers, pg 4 bottom of left column).

17. As per Claim 33, Wang and Lee disclose(s) the method of claim 28, wherein the plurality of characteristics includes a description of the plurality of array partitions (*Wang*, col 11 *Ins* 51-61) and said determining that the data access command pertains to the

Art Unit: 2185

array partition is based at least in part on the type of the RA group and the description of the plurality of array partitions (*as combined above, Wang, col 12 Ins 3-9*).

18. As per **Claim 36**, Wang and Lee disclose(s) a machine-accessible storage medium having instructions, which, when executed, results in the machine:

establishing a root partition (*Lee, location of the information on the Petal server, section 2 par 3*) on a storage medium, the root partition defining a plurality of characteristics (*Wang, configuration information, including LUN and ID information, RAID type, what multicast group they are in, what extent configurations are on the disk, stripe size, etc, col 11 Ins 51-61; Lee, information describing the current members of the storage system and the currently supported virtual disks, section 2 par 3*) of an array group (*Wang, array of disks in a logical group, col 8 Ins 31-41*) that includes a plurality of array partitions (*Wang, contents of disks in the logical group, col 8 Ins 31-41*), the plurality of characteristics including a type of the array group , which indicates how data is distributed across the plurality of array partitions (*Wang, RAID type, col 11 Ins 51-61*), or a parity rule of the array group,

establishing an array partition on the storage medium, the array partition being one of the plurality of array partitions (*Wang, the NetSCSI autoconfiguration protocol is used to assign LUN and ID information to each disk, col 11 Ins 51-61*),

receiving, via the network interface, a data access command multicast to the plurality of array partitions (*Wang, multicast the data to the disks, col 12 Ins 3-9*), and

determining that the data access command pertains to the array partition based at least in part on the plurality of characteristics (*Wang, the disks know*

Art Unit: 2185

*what LUNs they contain, and therefore what data would pertain to them, col 11
lns 51-61, col 12 lns 3-9).*

19. As per Claim 40, Wang and Lee disclose(s) an apparatus (*Lee, figure 2*) comprising:

a network interface (*scalable network, figure 2*); and
a controller (*Lee, user-level utilities operating on a remote machine; RPC
is known in the art as "Remote Procedure Call", implying the user-level utilities
are executed externally to Petal , section 3 par 3*) configured,

to transmit, via the network interface, a first partition command to
establish a root partition on a storage medium (*Lee, Petal RPC is used to
perform tasks such as adding a server, section 3 par 3*);

to transmit, via the network interface, a plurality of characteristics of an
array group (*Wang, array of disks in a logical group, col 8 lns 31-41*) which
includes a plurality of array partitions (*Wang, contents of disks in the logical
group, col 8 lns 31-41*), to be stored in the root partition (*Wang, configuration
information, including LUN and ID information, RAID type, what multicast group
they are in, what extent configurations are on the disk, stripe size, etc, col 11 lns
51-61; Lee, information describing the current members of the storage system
and the currently supported virtual disks, section 2 par 3*);

to transmit, via the network interface, a second partition command to
establish an array partition of the plurality of array partitions, on the storage
medium (*Wang, the NetSCSI autoconfiguration protocol is used to assign LUN
and ID information to each disk, col 11 lns 51-61; Lee, occasionally a virtual
disk's redundancy scheme over the set of servers over which it is mapped is
changed, section 2.3*); and

to multicast, via the network interface, a packet to the plurality of array partitions (*Wang, multicast the data to the disks, col 12 lns 3-9*), the packet having a data access command and a logical block address (LBA) to which the data access command pertains (*Wang, col 20 lns 31-42*), the LBA associated with only a subset of the plurality of array partitions (*Wang, the disks know what data they contain, and therefore what data would pertain to them, col 11 lns 51-61, col 12 lns 3-9*).

20. As per **Claim 41**, Wang and Lee disclose(s) the apparatus of claim 40, wherein the controller is further configured to transmit a partition command to each of a plurality of storage appliances associated with a respective plurality of storage media to establish the plurality of array partitions (*Wang, each disk is configured, col 11 lns 51-61*).

21. As per **Claim 44**, Wang and Lee disclose(s) the method of claim 28, further comprising: receiving, via the network interface, one or more commands from a host; and providing the root partition and the array partition based at least in part on the received one or more commands (*Lee, user-level utilities operating on a remote machine; RPC is known in the art as "Remote Procedure Call", implying the user-level utilities are executed externally to Petal, section 3 par 3*).

22. As per **Claim 45**, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the plurality of characteristics includes the type of the array group (*Wang, configuration includes RAID type, col 11 lns 51-61*), and the type comprises a RAID type 0, 1, 4, or 5 (*Wang, col 4 lns 18-29*).

23. As per **Claim 46**, Wang and Lee disclose(s) the storage appliance of claim 17, wherein the controller is further configured: to buffer data transmitted in the data access command (*Wang, col 7 lns 15-17*); to transfer, via the network interface based at least in part on the data access command, data from the array partition to a parity partition of the

Art Unit: 2185

plurality of array partitions (*Wang, col 7 lns 11-17*); and to save the buffered data in the array partition (*Wang, col 7 lns 15-17*).

Claim Rejections - 35 USC ' 103 – Wang, Lee and Vigue

24. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

25. **Claims 22 and 32** are rejected under 35 U.S.C. 103(a) as being unpatentable over Wang et al. (*US Patent 6,834,326*) and Lee et al. ("*Petal: Distributed Virtual Disks*") as combined above with respect to Claim 28, in further view of Vigue et al. (*US Patent 6,983,326*).

26. As per **Claims 22 and 32**, Wang and Lee disclose(s) the storage appliance of claim 17, but for the purposes of this rejection does not disclose Claim 22.

Vigue discloses a controller configured to receive, via a network interface, a command (*receive "I Need" packer, fig 3*); to receive, via the network interface, a response to the command (*receive "I found" packet, fig 3*); and to disregard the command based at least in part on the response (*canceled request, fig 3*).

Wang as combined with Lee and Vigue are analogous art in that they deal with peer network devices. At the time of the invention, it would have been obvious to a person having ordinary skill in the art to modify Wang and Lee to ignore a request if it has already been fulfilled, per the teachings of Vigue. The motivation would have been obvious to a person having ordinary skill in the art, and would have been that it would reduce network traffic. Therefore, it would have been obvious to modify Wang and Lee

Art Unit: 2185

per the teachings of Vigue for the benefit of reduced network traffic, to obtain the invention of Claims 22 and 32.

III. CLOSING COMMENTS

27. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office Action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP ' 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

a. STATUS OF CLAIMS IN THE APPLICATION

28. The following is a summary of the treatment and status of all claims in the application as recommended by M.P.E.P. ' 707.07(i):

a(1). CLAIMS NO LONGER IN THE APPLICATION

29. Claims 1-16, 29 and 42 were cancelled by amendment.

a(2). CLAIMS REJECTED IN THE APPLICATION

30. Per the instant office action, Claims 17-28, 30-41 and 43-46 have received an action on the merits and are subject of a final action.

Art Unit: 2185

b. DIRECTION OF FUTURE CORRESPONDENCES

31. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sam Dillon whose telephone number is 571- 272-8010.

The examiner can normally be reached on 9:30-6:00.

32. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Sanjiv Shah can be reached on 571-272-4098. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

IMPORTANT NOTE

33. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Hong Kim/
Primary Examiner, Art Unit 2185

Sam Dillon
Examiner
Art Unit 2185